Coherence ethics revisited, formalized

Benjamin James

November 17, 2025

Coherence ethics needs to be revisited beyond its original formulation into a unified account of how moral systems evolve, stabilize, and act within complex environments. My initial framework began as the self-contained development of a model of ethics grounded not in rules, outcomes, or cultural norms, but in the behavior of adaptive systems seeking to persist under entropy. The purpose here is not to revise the earlier work but to clarify its implications and demonstrate its independence as a stand-alone theory, while tightening its relationship to multi-scale stability.

At the core of the framework lies a simple claim; *ethical action is the behavior that increases coherence across relevant scales without collapsing option-space or externalizing instability*. This principle, though minimal, has wide reach. It implies moral norms emerge because coherent systems outperform incoherent ones, responsibility increases with an agent's capacity to model recursive effects, and that many classical moral dilemmas become artifacts of incomplete scale-tracking rather than deep philosophical contradictions. Ethics is not a set of prescriptions but a dynamic property of systems attempting to survive, learn, and coordinate.

The idea that morality is coherence-maximizing recontextualizes long-standing debates. The is–ought divide dissolves because norms are not arbitrary impositions but stability gradients. Systems that behave in ways aligned with coherence ascent persist, while those that do not fail. Moral relativism becomes structurally constrained rather than absolute; cultures vary, but only some of their practices scale because only those practices contribute positively to long-term coherence. Deontological rules reveal themselves as low-entropy heuristics useful in predictable environments, while consequentialist reasoning becomes essential in high-uncertainty or rapidly shifting contexts. Neither is correct or sufficient alone; each describes different regions of the same landscape.

This perspective also clarifies why humans experience moral intuition as both rigid and flexible. Individuals rely on familiar rules to reduce cognitive load, but when environments destabilize, the mind naturally shifts into outcome-evaluating modes. The same structure applies to legal systems, institutions, and societies. As long as conditions remain stable, rules accumulate, codify, and harden. When coherence begins to erode through technological disruption, ecological stress, or political fragmentation, those rules reveal their brittleness. Coherence ethics explains this transition

as an adaptive necessity; a system too attached to static rules cannot respond to entropic pressure, while a system that abandons rules entirely collapses into noise.

Under this lens, moral responsibility expands with the depth of recursion an agent can meaningfully evaluate. A child, animal, AI, and an advanced self-modeling intelligence occupy different tiers of ethical capacity because each can project and regulate coherence to different degrees. This structure avoids anthropocentrism while preserving meaningful distinctions among agents. It also offers a clear criterion for when artificial systems should be treated as moral subjects. Not when they mimic human behavior or language, but when they demonstrate the ability to refine their own coherence across time and scale.

This provides a new interpretation of well-known moral tensions. Consider trolley-style dilemmas. Traditional analysis frames them as conflicts between utilitarian and deontological reasoning. Coherence Ethics instead treats them as disruptions in trust and predictability. Pulling a lever constitutes a minor intervention with limited relational fallout; pushing a person introduces a deeply destabilizing violation of embodied trust and social inference. The difference is not mystical or sentimental; it reflects scales of coherence disruption. Similarly, altruism, often treated as mysterious or irrational, becomes legible as a stabilizing force that sacrifices local coherence for global benefit. Justice and mercy, often presented as moral opposites, become complementary tools for stabilizing systems at different scales. Justice enforces macro-level predictability, while mercy prevents micro-level overcorrection.

Ethical reasoning under this model follows a coherent arc. One first identifies which scales are active and which are at risk, such as the individual, the group, the ecological substrate, the institutional framework, or the long arc of civilization itself. One analyzes how current rules and habits contribute to or erode stability. One selects the action that best increases overall coherence without inducing unnecessary volatility. And crucially, one revises the model after each decision, allowing the system to self-correct over time. Ethics becomes an ongoing calibration process rather than a fixed code.

Taken as a whole, coherence ethics offers a substrate capable of integrating moral reasoning across domains that typically resist unification, such as personal integrity, collective governance, ecological stewardship, artificial intelligence alignment, and the long-term trajectory of intelligent systems. It neither prescribes specific behaviors nor collapses into relativism. Instead, it supplies the structural logic by which any viable ethical system must operate if it hopes to scale across agents, cultures, technologies, and epochs.

By grounding morality in coherence, the measurable, structural condition that enables intelligent systems to persist and coordinate, it becomes possible to resolve long-standing philosophical debates, avoid brittle doctrines, and build ethical architectures capable of withstanding the accelerating pressures of modern and posthuman futures. Coherence ethics, when articulated as a standalone theory, is not merely one ethical option among many. It is the minimal and necessary condition for any system that seeks to endure and evolve within an entropic universe.

Axioms and Mathematical Structure

Axiom 1 – Coherence Persistence

A system behaves ethically when coherence increases faster than it decays,

$$\frac{dC}{dt} = \Gamma C^n - \lambda C + \eta(\nabla S \cdot \nabla \Omega)$$

with the ethical persistence condition,

$$\frac{dC}{dt} > 0$$

and where,

C = coherence density

Γ = stabilization gain

n = recursion nonlinearity

λ = entropy leakage

η = responsiveness

∇S = entropy gradient

$\nabla \Omega$ = option-space gradient

Axiom 2 – Ethical Trajectory

Ethical action follows entropy descent and coherence ascent,

$$\frac{dt}{ds} = -\nabla S(s) + \nabla C(s)$$

An action is ethical if it moves the system toward higher coherence and lower entropy.

Axiom 3 – Information Value

Only information that increases coherence is morally relevant,

$$IC = \sum_i p_i \, \mu(\mathrm{C}_i)$$

Where $\mu(\mathrm{C}_i)$ is the coherence-yield of information state i.

Axiom 4 – Recursive Arbitration

When multiple actions are possible, the ethical choice maximizes long-term coherence curvature while preserving option-space,

$$A^* = \arg\max_i \int_\tau \left(\frac{\mathrm{d}^2\mathrm{C}_i}{\mathrm{dt}^2}\right) \mu\left(\frac{\partial \mathrm{C}_i}{\partial \mathrm{C}_{\mathrm{system}}}\right) \Omega_i^{\alpha} d\tau$$

and where,

$\frac{d^2 C_i}{dt^2}$ = coherence curvature

$\mu\left(\frac{\partial \mathrm{C}_i}{\partial \mathrm{C}_{\mathrm{system}}}\right)$ = coupling term

Ω_i = option-space retention

α = long-horizon weighting exponent

Axiom 5 – Pruning Rule

A structure must be removed when it decreases coherence and expands destabilizing variability,

$$\mathrm{Prune}(S_i) \Leftrightarrow \Delta C_i < 0 \wedge \frac{d\Omega}{dt} > 0$$

Cross-Scale Coherence

Ethical evaluation aggregates coherence across scales,

$$C_{\mathrm{total}}(x) = \sum_s w_s(x) \, C_s(x)$$

weights w_s increase for scales nearest to instability, and ethical action is,

$$x^* = \arg\max_x C_{\mathrm{total}}(x)$$

Responsibility Function

Moral responsibility grows with recursive modeling depth,

$$R_{\text{agent}} = k \cdot D_{\text{agent}}$$

Where D_{agent} is recursion depth.

Ethical Collapse Conditions

Ethics fails when either of the following holds,

1. $\Delta C < 0$ for horizon τ
2. $\frac{d\Omega}{dt} < 0$ due to overconstraint

Decision Operator

Ethical decision-making combines coherence gradient, rule stability, and outcome stability,

$$E(x) = \nabla C(x) + k_1 R(x) + k_2 O(x)$$

and ethically optimal action,

$$x^* = \arg\max_x E(x)$$

These axioms and equations provide the necessary and sufficient structure for coherence ethics to function as a mathematically grounded, recursively stable ethical system capable of scaling across biological, social, artificial, and civilizational domains.

References

James, B. (2025). Coherence Ethics: A Universal Framework for Adaptive Moral Reasoning. PhilPapers. https://philpapers.org/rec/JAMCEA-6

James, B. (2025). Formal Foundations of Adaptive Coherence: A Recursive Metric of Reality. PhilPapers. https://philpapers.org/rec/JAMFFO